

Divergence Loss and the Stein Phenomenon

Malay Ghosh
University of Florida

June 19, 2019

Outline

- Introduction
- Squared Error Loss
- Divergence Loss
- Inadmissibility Results
- Stein's Harmonic Prior
- A Two Sample Problem
- Final Remarks

Introduction

- The natural estimator of the normal mean, the sample mean, is its UMVUE, the MLE, and also the best equivariant estimator under translations of the sample space.
- This estimator is minimax under a very general class of losses in any arbitrary dimension.
- Blyth (1951) proved the admissibility of the sample mean in the normal case for a very general class of losses, including but not limited to the squared error loss, in one dimension.
- Stein (1956) made the surprising discovery that although the sample mean is an admissible estimator of the normal mean in one or two dimensions under squared error loss, it is not so in three or higher dimensions under the same loss.

- Baranchik (1970), under squared error loss, provided a very general class of minimax estimators including the James-Stein estimator which dominates the sample mean in three or higher dimensions.
- A very important subclass of the Baranchik class of estimators, namely proper Bayes minimax estimators under certain hierarchical priors, was found by Strawderman (1971).
- Efron and Morris, in a series of articles, provided an empirical Bayes (EB) interpretation of the James-Stein estimator.
- Extending their ideas further, the entire Baranchik class of estimators, including those of Strawderman, can be given simple EB interpretation.

- James-Stein estimators or their variants which dominate the sample mean continue to do so when it comes to prediction of a future observation from the normal distribution with the same mean under squared error loss.
- This is a consequence of a more general result which establishes duality between estimation and prediction under squared error loss.
- Robert (2001) points out that often it is natural to use losses which compare directly the densities $f(\cdot|\mathbf{x})$ and $f(\cdot|\mathbf{a})$, where θ is the true parameter.
- Robert refers to such losses as "intrinsic losses".

- The two most well-used divergence measures between two densities are the Kullback-Leibler (KL) or the entropy distance and the Bhattacharyya-Hellinger (BH) distance (Hellinger, 1909; Bhattacharyya, 1943).
- However, if $\mathbf{X} = (X_1, \dots, X_p)^T \sim N(\boldsymbol{\theta}, v_x \mathbf{I}_p)$, the KL distance between $f(\cdot|\boldsymbol{\theta})$ and $f(\cdot|\mathbf{a})$ is

$$E_{\boldsymbol{\theta}} \left[\log \frac{f(\mathbf{X}|\boldsymbol{\theta})}{f(\mathbf{X}|\mathbf{a})} \right] = \frac{1}{2v_x} \|\boldsymbol{\theta} - \mathbf{a}\|^2.$$

- Hence, point estimation based on the KL loss is tantamount to squared error loss for known v_x .

- KL and BH losses are special cases of a more general divergence loss considered in other contexts (Renyi, 1961; Amari, 1982; Cressie and Read, 1984).
- This loss is given by

$$L_{\beta}(\theta, \mathbf{a}) = \frac{1 - \int f^{1-\beta}(\mathbf{x}|\theta)f^{\beta}(\mathbf{x}|\mathbf{a})d\mathbf{x}}{\beta(1 - \beta)}$$

- The above loss is to be interpreted as its limit when $\beta \rightarrow 0$ or $\beta \rightarrow 1$.

- $\lim_{\beta \rightarrow 0} L_{\beta}(\theta, \mathbf{a}) = E_{\theta}[\log \frac{f_{\theta}(\mathbf{X})}{f_{\mathbf{a}}(\mathbf{X})}]$.

- $\lim_{\beta \rightarrow 1} L_{\beta}(\theta, \mathbf{a}) = E_{\mathbf{a}}[\log \frac{f_{\mathbf{a}}(\mathbf{X})}{f_{\theta}(\mathbf{X})}]$.



$$\begin{aligned} L_{1/2}(\theta, \mathbf{a}) &= 4[1 - \int f^{1/2}(\mathbf{x}|\theta)f^{1/2}(\mathbf{x}|\mathbf{a})d\mathbf{x}] \\ &= 2 \int [f^{1/2}(\mathbf{x}|\theta) - f^{1/2}(\mathbf{x}|\mathbf{a})]^2 d\mathbf{x}. \end{aligned}$$

- Recently Komaki (2001) and George et. al (2006) have considered improved minimax predictive densities under the KL loss.
- They have developed various shrinkage versions of predictive densities which dominate under this loss the Bayes predictive density with the uniform prior for a future observation conditionally independent of the sampled observations.
- George et. al (2006) have also explored various interesting duality results between multivariate estimation and prediction in the normal problem under this loss.

- In contrast, Ghosh, Mergel and Datta (2006) considered the general divergence loss, established duality between estimation and prediction, and showed that the general Baranchik class of estimators continues to dominate the sample mean under this general divergence loss.
- Moreover, the EB interpretation of the Baranchik class of estimators continues to hold under this loss.
- The key feature of this talk is to explore this duality between estimation and prediction under the general divergence loss.

Squared Error Loss

- Let $X_1, \dots, X_p | \theta_1, \dots, \theta_p$ be independent $N(\theta_i, v_x)$, $i = 1, \dots, p$, $v_x > 0$ known.
- We write $\mathbf{X} = (X_1, \dots, X_p)^T$ and $\boldsymbol{\theta} = (\theta_1, \dots, \theta_p)^T$.
- For $p \geq 3$, under the loss $L(\boldsymbol{\theta}, \mathbf{a}) = \|\mathbf{a} - \boldsymbol{\theta}\|^2$, $\delta^{\text{JS}}(\mathbf{X}) = \left[1 - \frac{(p-2)v_x}{\|\mathbf{X}\|^2}\right]\mathbf{X}$ dominates \mathbf{X} .
- Let $S = \|\mathbf{X}\|^2/v_x$. Baranchik (1970) showed that the general class of estimators $\delta^{\text{BA}}(\mathbf{X}) = \left[1 - \frac{\tau(S)}{S}\right]\mathbf{X}$ dominates \mathbf{X} if (i) $\tau(S)$ is non-decreasing in S and (ii) $0 < \tau(S) < 2(p-2)$.
- If $\tau(S) = c$, $0 < c < 2(p-2)$, the optimal $c = p-2$, which gives the James-Stein estimator.

- The general Baranchik class of estimators can be given an EB interpretation.
- $\mathbf{X}|\boldsymbol{\theta} \sim N(\boldsymbol{\theta}, v_x \mathbf{I}_p)$ and $\boldsymbol{\theta} \sim N(\mathbf{0}, A v_x \mathbf{I}_p)$, v_x known, A unknown.
- The Bayes estimator of $\boldsymbol{\theta}$ under squared error loss is $(1 - B)\mathbf{X}$, where $B = (1 + A)^{-1}$.
- To estimate B , note that marginally $\mathbf{X} \sim N(\mathbf{0}, B^{-1} v_x \mathbf{I}_p)$.
- $S = \|\mathbf{X}\|^2 / v_x$ is complete sufficient for B .
- A general class of EB estimators of B is of the form $(1 - \hat{B}(S))\mathbf{X}$.
- In particular, since marginally, $S \sim B^{-1} \chi_p^2$, the UMVUE of B is $(p - 2)/S$.
- The resulting EB estimator of B is $(p - 2)/S$ leading to the James-Stein estimator $(1 - \frac{p-2}{S})\mathbf{X}$.

- An important class of Bayes estimators dominating \mathbf{X} was provided by Stein (1981).
- Consider a general prior $\pi(\boldsymbol{\theta})$ of $\boldsymbol{\theta}$.
- Let $m_\pi(\mathbf{X}) = \int \phi(\mathbf{X} | (\boldsymbol{\theta}, v_x \mathbf{I}_p)) \pi(\boldsymbol{\theta}) d\boldsymbol{\theta}$ be finite for all \mathbf{X} .
- Then the Bayes estimator of $\boldsymbol{\theta}$ is $\delta_\pi(\mathbf{X}) = \mathbf{X} + v_x \nabla \log m_\pi(\mathbf{X})$, where ∇ denotes the gradient vector.
- Stein showed that if either m_π or $m_\pi^{1/2}$ is superharmonic, i.e. it has a negative Laplacian, then $\delta_\pi(\mathbf{X})$ dominates \mathbf{X} .
- In particular, Stein proposed the harmonic prior $\pi(\boldsymbol{\theta}) = \|\boldsymbol{\theta}\|^{-(p-2)}$.

- Duality between dominance under squared error loss over the estimator \mathbf{X} of θ and the predictor \mathbf{Y} where $\mathbf{X} \perp \mathbf{Y} | \theta$ and $E(\mathbf{X}) = E(\mathbf{Y}) = \theta$.
- Suppose an estimator $e(\mathbf{X})$ of θ dominates \mathbf{X} under squared error loss. Then
-

$$\begin{aligned}
 E_{\theta} \|e(\mathbf{X}) - \mathbf{Y}\|^2 & \\
 &= E_{\theta} \|\{e(\mathbf{X}) - \theta\} - \{\mathbf{Y} - \theta\}\|^2 \\
 &= E_{\theta} \|e(\mathbf{X}) - \theta\|^2 + E_{\theta} \|\mathbf{Y} - \theta\|^2 \\
 &< E_{\theta} \|\mathbf{X} - \theta\|^2 + E_{\theta} \|\mathbf{Y} - \theta\|^2 \\
 &= E_{\theta} \|\mathbf{X} - \mathbf{Y}\|^2.
 \end{aligned}$$

- Conversely, suppose $E_{\theta} \|e(\mathbf{X}) - \mathbf{Y}\|^2 < E_{\theta} \|\mathbf{X} - \mathbf{Y}\|^2$.
- Then subtracting $E_{\theta} \|\mathbf{Y} - \theta\|^2$ from both sides, $E_{\theta} \|e(\mathbf{X}) - \theta\|^2 < E_{\theta} \|\mathbf{X} - \theta\|^2$.

Divergence Loss

- Let $\mathbf{X} \perp \mathbf{Y}|\theta$ with corresponding pdf's $p(\mathbf{x}|\theta)$ and $p(\mathbf{y}|\theta)$.
- Recall the divergence loss $L_\beta(\theta, \mathbf{a}) = \frac{1 - \int f^{1-\beta}(\mathbf{x}|\theta) f^\beta(\mathbf{x}|\mathbf{a}) d\mathbf{x}}{\beta(1-\beta)}$.
- Let $\pi(\theta|\mathbf{x})$ denote the posterior pdf of θ given \mathbf{x} .
- Lemma 1. Under the divergence loss and the prior π , the Bayes predictive density of \mathbf{Y} is $\pi_D(\mathbf{y}|\mathbf{x}) = k^{1/(1-\beta)}(\mathbf{y}, \mathbf{x}) / \int k^{1/(1-\beta)}(\mathbf{y}, \mathbf{x}) d\mathbf{y}$, where $k(\mathbf{y}, \mathbf{x}) = \int p^{1-\beta}(\mathbf{y}|\theta) \pi(\theta|\mathbf{x}) d\theta$.
- In the special case when $\beta \rightarrow 0$, namely the KL loss, $\pi_{KL}(\mathbf{y}|\mathbf{x}) = \int p(\mathbf{y}|\theta) p(\theta|\mathbf{x}) d\theta$, the usual predictive pdf (Atchison, 1975).

- A useful result.
- Suppose \mathbf{Z}_1 and \mathbf{Z}_2 are two independent p -component random vectors with pdf's f_1 and f_2 and distribution functions F_1 and F_2 , respectively.
- Let h and H denote respectively the pdf and distribution function of $\mathbf{Z}_1 - \mathbf{Z}_2$.
- Then $h(\mathbf{0}) = \int f_1(\mathbf{z})f_2(\mathbf{z})d\mathbf{z}$.
- The proof follows by first observing $H(\mathbf{x}) = \int F_1(\mathbf{x} + \mathbf{z})f_2(\mathbf{z})d\mathbf{z}$.
- Differentiating both sides with respect to \mathbf{x} , one gets $h(\mathbf{x}) = \int f_1(\mathbf{x} + \mathbf{z})f_2(\mathbf{z})d\mathbf{z}$, which leads to the value of $h(\mathbf{0})$.
- In particular,
$$\int \phi(\mathbf{z}|\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)\phi(\mathbf{z}|\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2)d\mathbf{z} = \phi(\mathbf{0}|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2, \boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2).$$
- Also, $L_\beta(\boldsymbol{\theta}, \mathbf{a}) = \exp\left[-\frac{\beta(1-\beta)}{2v_x}\|\boldsymbol{\theta} - \mathbf{a}\|^2\right]$.

- Consider once again $\mathbf{Y} \perp \mathbf{X} | \boldsymbol{\theta}$ with $\mathbf{Y} \sim N(\boldsymbol{\theta}, v_y \mathbf{I}_p)$ and $\mathbf{X} \sim N(\boldsymbol{\theta}, v_x \mathbf{I}_p)$, $v_x (> 0)$ and $v_y (> 0)$ are known.
- $\boldsymbol{\theta}$ has the $N(\boldsymbol{\mu}, A \mathbf{I}_p)$ prior. $(1 - B)\mathbf{X} + B\boldsymbol{\mu}$, where $B = v_x / (v_x + A)$.
- The goal here is to maximize

$$\int \exp\left\{-\frac{\beta(1-\beta)}{2v_x} \|\boldsymbol{\theta} - \mathbf{a}\|^2\right\} \phi(\boldsymbol{\theta} | (1-B)\mathbf{X} + B\boldsymbol{\mu}, v_x(1-B)\mathbf{I}_p) d\boldsymbol{\theta}$$
 with respect to \mathbf{a} .
- This integral is proportional to $\exp\left[-\frac{\|\mathbf{a} - (1-B)\mathbf{X} - B\boldsymbol{\mu}\|^2}{2v_x\{\beta^{-1}(1-\beta)^{-1} + 1 - B\}}\right]$.
- This is maximized when $\mathbf{a} = (1 - B)\mathbf{X} + B\boldsymbol{\mu}$.
- Thus the Bayes estimator of $\boldsymbol{\theta}$ is $(1 - B)\mathbf{X} + B\boldsymbol{\mu}$, where $B = v_x / (v_x + A)$, same as squared error loss.

- For prediction, writing $(1 - B)\mathbf{X} + B\boldsymbol{\mu} = \hat{\boldsymbol{\theta}}_B(\mathbf{X})$,

$$[\hat{p}_{\mathbf{Y}}(\mathbf{y}|\mathbf{X})]^{1-\beta} \propto \left[\int \phi^{1-\beta}(\mathbf{y}|\boldsymbol{\theta}, v_y \mathbf{I}_p) \phi(\boldsymbol{\theta}|\hat{\boldsymbol{\theta}}_B(\mathbf{X}), v_x(1 - B)\mathbf{I}_p) d\boldsymbol{\theta} \right]$$

$$\propto [\phi(\mathbf{y}|(1 - B)\mathbf{X} + B\boldsymbol{\mu}, \{v_x(1 - B)(1 - \beta) + v_y\}\mathbf{I}_p)]^{1-\beta}.$$

- If $\pi(\boldsymbol{\theta}) = 1$, $B \rightarrow 0$.
- Then the Bayes estimator of $\boldsymbol{\theta}$ is \mathbf{X} .
- The Bayes predictive density of \mathbf{Y} is $\phi(\mathbf{y}|\mathbf{X}, (v_x(1 - \beta) + v_y)\mathbf{I}_p)$.
- We denote by δ_0 the plug-in predictive density $\hat{p}_{pl}(\mathbf{y}|\mathbf{X}) = \phi(\mathbf{y}|\mathbf{X}, v_y \mathbf{I}_p)$.

- For estimating the normal mean, the Bayes estimator of θ under the prior $\pi(\theta) = 1$ is \mathbf{X} for the general divergence loss as well as squared error loss.
- Similarly, under the divergence loss, the Bayes predictive density of $\phi(\mathbf{y}|\theta, v_y I_p)$ under the prior $\pi(\theta) = 1$ is $\hat{p}_U(\mathbf{y}|\mathbf{X}) = \phi(\mathbf{y}|\mathbf{X}, (v_x + v_y)I_p)$.
- $\phi(\mathbf{y}|\mathbf{X}, (v_x + v_y)I_p)$ has smaller risk than that of the plug-in predictor $\phi(\mathbf{y}|\mathbf{X}, v_y I_p)$.
- \mathbf{X} is a minimax estimator of θ under the general divergence loss in any arbitrary dimension.
- Similarly, $\hat{p}_U(\mathbf{y}|\mathbf{X}) = \phi(\mathbf{y}|\mathbf{X}, (v_x + v_y)I_p)$ is a minimax predictor of $\phi(\mathbf{y}|\theta, v_y I_p)$ under the general divergence loss in any arbitrary dimension.
- While Stein type shrinkage estimators are targeted to dominate \mathbf{X} , the goal is to improve on $\hat{p}_U(\mathbf{y}|\mathbf{X})$ via shrinkage.
- This is what one wants to accomplish under the general divergence loss.

Inadmissibility Results

- $\hat{p}_U(\mathbf{y}|\mathbf{X})$ is an admissible predictor of $\phi(\mathbf{y}|\boldsymbol{\theta}, \mathbf{I}_p)$ for $p = 1$.
- Its admissibility is an open question for $p = 2$. (?)
- $\hat{p}_U(\mathbf{y}|\mathbf{X})$ is an inadmissible predictor of $p(\mathbf{y}|\boldsymbol{\theta}, v_y \mathbf{I}_p)$ for $p \geq 3$.
- The duality between estimation and prediction can be seen via the Baranchik class of estimators and predictors.
- Let $S = \|\mathbf{X}\|^2/v_x$.
- The Baranchik class of estimators for $\boldsymbol{\theta}$ is given by $\boldsymbol{\delta}^\tau(\mathbf{X}) = \{1 - \tau(S)/S\}\mathbf{X}$.
- Analogously, the Baranchik class of predictors for $\phi(\mathbf{y}|\boldsymbol{\theta}, v_y \mathbf{I}_p)$ is given by $\phi(\mathbf{y}|\boldsymbol{\delta}^\tau(\mathbf{X}), ((1 - \beta)v_x + v_y)\mathbf{I}_p)$.

- For the estimation problem,

$$L(\boldsymbol{\theta}, \boldsymbol{\delta}^\tau(\mathbf{X})) = \frac{1}{\beta(1-\beta)} \left[1 - \exp \left\{ -\frac{\beta(1-\beta)}{2v_x} \|\boldsymbol{\delta}^\tau(\mathbf{X}) - \boldsymbol{\theta}\|^2 \right\} \right].$$

- For the prediction problem,

$$\begin{aligned} & L(\phi(\mathbf{y}|\boldsymbol{\theta}, \mathbf{I}_p), \phi(\mathbf{y}|\boldsymbol{\delta}^\tau(\mathbf{X}), \{(1-\beta)v_x + v_y\}\mathbf{I}_p\}) \\ &= \frac{1}{\beta(1-\beta)} \left[1 - \frac{(v_y)^{p\beta/2} \{(1-\beta)v_x + v_y\}^{p(1-\beta)/2}}{\{(1-\beta)^2 v_x + v_y\}^{p/2}} \exp \left\{ -\frac{\beta(1-\beta) \|\boldsymbol{\delta}^\tau(\mathbf{X}) - \boldsymbol{\theta}\|^2}{2\{(1-\beta)^2 v_x + v_y\}} \right\} \right]. \end{aligned}$$

- Theorem. Let $p \geq 3$. Assume (i) $0 < \tau(t) < 2(p-2)$ and (ii) $\tau(t)$ is a differentiable nondecreasing function of t . Then
 - (a) $R(\boldsymbol{\theta}, \boldsymbol{\delta}^\tau(\mathbf{X})) < R(\boldsymbol{\theta}, \mathbf{X})$ and
 - (b) $R(\phi(\mathbf{y}|\boldsymbol{\theta}, v_y \mathbf{I}_p), \phi(\mathbf{y}|\boldsymbol{\delta}^\tau(\mathbf{X}), ((1-\beta)v_x + v_y)\mathbf{I}_p))$
 $< R(\phi(\mathbf{y}|\boldsymbol{\theta}, v_y \mathbf{I}_p), \phi(\mathbf{y}|\mathbf{X}, ((1-\beta)v_x + v_y)\mathbf{I}_p))$.

- The above result can immediately be extended to shrinkage towards an arbitrary regression surface.
- Suppose now that $\mathbf{X}|\boldsymbol{\theta} \sim \mathcal{N}(\boldsymbol{\theta}, v_x \mathbf{I}_p)$.
- $\boldsymbol{\theta} \sim \mathcal{N}(\mathbf{K}\boldsymbol{\beta}, A\mathbf{I}_p)$ where \mathbf{K} is a known $p \times r$ matrix of rank $r (< p)$ and $\boldsymbol{\beta}$ is $r \times 1$ regression coefficient.
- Define $\mathbf{P} = \mathbf{K}(\mathbf{K}^T \mathbf{K})^{-1} \mathbf{K}^T$ and $R^* = \|\mathbf{X} - \mathbf{P}\mathbf{X}\|^2 / v_x$.
- Let $p \geq r + 3$. Assume (i) $0 < \tau(t) < 2(p - 2)$ and (ii) $\tau(t)$ is a differentiable nondecreasing function of t . Then $\mathbf{X} - \{\tau(R^*)/R^*\}(\mathbf{X} - \mathbf{P}\mathbf{X})$ dominates \mathbf{X} under the general divergence loss.
- A similar dominance result holds for prediction of $\phi(\mathbf{y}|\boldsymbol{\theta}, v_y \mathbf{I}_p)$.

- Consider now the case $\mathbf{X} \sim N(\boldsymbol{\theta}, v_x \mathbf{I}_p)$, where both $\boldsymbol{\theta}$ and v_x are unknown.
- $S \sim \frac{v_x}{m+2} \chi_m^2$ independent of \mathbf{X} .
- This is the typical balanced one way ANOVA model where \mathbf{X} is the vector of treatment means and S is a multiple of the error mean square.
- The goal continues to be estimation of $\boldsymbol{\theta}$ under the general divergence loss $L(\boldsymbol{\theta}, \mathbf{a}) = \frac{1 - \exp\left[-\frac{\beta(1-\beta)}{2v_x} \|\boldsymbol{\theta} - \mathbf{a}\|^2\right]}{\beta(1-\beta)}$, $0 < \beta < 1$.
- Noting that $\|\mathbf{X} - \boldsymbol{\theta}\|^2 / v_x \sim \chi_p^2$, the risk $R(\boldsymbol{\theta}, \mathbf{X}) = \frac{1 - [1 + \beta(1-\beta)]^{-p/2}}{\beta(1-\beta)}$.

- Following Baranchik (1970) and Efron and Morris (1976), we consider the rival set of estimators $\delta(\mathbf{X}) = \left(1 - \frac{\tau(\|\mathbf{X}\|^2/S)}{\|\mathbf{X}\|^2/S}\right) \mathbf{X}$.

- Consider the loss $L(\boldsymbol{\theta}, \delta(\mathbf{X})) = \frac{1 - \exp\left[-\frac{\beta(1-\beta)}{2v_x} \|\delta(\mathbf{X}) - \boldsymbol{\theta}\|^2\right]}{\beta(1-\beta)}$.

- Theorem. Let $p \geq 3$. Assume
 - (i) $0 < \tau(t) < 2(p-2)$ for all $t > 0$;
 - (ii) $\tau(t)$ is a differentiable nondecreasing function of t for $t > 0$.

Then $R(\boldsymbol{\theta}, \delta(\mathbf{X})) < R(\boldsymbol{\theta}, \mathbf{X})$ for all $\boldsymbol{\theta} \in R^p$.

- Finally, let $\mathbf{Z}_1, \dots, \mathbf{Z}_n$ ($n \geq 2$) be iid $N_p(\boldsymbol{\theta}, \boldsymbol{\Sigma})$, where both $\boldsymbol{\theta}$ and $\boldsymbol{\Sigma}$ are unknown.
- Let $\mathbf{X} = \bar{\mathbf{Z}}$ and $\mathbf{S} = (n-1)^{-1} \sum_{i=1}^n (\mathbf{Z}_i - \bar{\mathbf{Z}})(\mathbf{Z}_i - \bar{\mathbf{Z}})^T$.
- The divergence loss is now given by

$$\begin{aligned}
 L(\boldsymbol{\theta}, \mathbf{a}) &= \frac{1 - \int \phi^{1-\beta}(\mathbf{x}|\boldsymbol{\theta}, \boldsymbol{\Sigma}/n) \phi^\beta(\mathbf{x}|\mathbf{a}, \boldsymbol{\Sigma}/n) d\mathbf{x}}{\beta(1-\beta)} \\
 &= \frac{1 - \exp\left[-\frac{n\beta(1-\beta)}{2}(\mathbf{a} - \boldsymbol{\theta})^T \boldsymbol{\Sigma}^{-1}(\mathbf{a} - \boldsymbol{\theta})\right]}{\beta(1-\beta)}.
 \end{aligned}$$

- Since $n(\mathbf{X} - \boldsymbol{\theta})^T \boldsymbol{\Sigma}^{-1}(\mathbf{X} - \boldsymbol{\theta}) \sim \chi_p^2$,

$$R(\boldsymbol{\theta}, \mathbf{X}) = \frac{[1 - \{1 + \beta(1-\beta)\}^{-p/2}]}{\beta(1-\beta)}.$$

- Consider now the general class of estimators

$$\delta(\mathbf{X}, \mathbf{S}) = \left[1 - \frac{\tau(n\mathbf{X}^T \mathbf{S}^{-1} \mathbf{X})}{n\mathbf{X}^T \mathbf{S}^{-1} \mathbf{X}} \right] \mathbf{X}.$$

- Consider the loss

$$L(\boldsymbol{\theta}, \delta(\mathbf{X}, \mathbf{S})) = \frac{1 - \exp\left[-\frac{n\beta(1-\beta)}{2}(\delta(\mathbf{X}, \mathbf{S}) - \boldsymbol{\theta})^T \boldsymbol{\Sigma}^{-1}(\delta(\mathbf{X}, \mathbf{S}) - \boldsymbol{\theta})\right]}{\beta(1-\beta)}.$$

- Theorem. Let $p \geq 3$. Assume
 - $0 < \tau(t) < 2(p-2)(n-1)/(n-p+2)$;
 - $\tau(t)$ is non-decreasing differentiable function of t .
 Then $R(\boldsymbol{\theta}, \delta(\mathbf{X}, \mathbf{S})) < R(\boldsymbol{\theta}, \mathbf{X})$.

Stein's Harmonic Prior

- Stein's harmonic prior yielded Bayes estimators of θ dominating \mathbf{X} when $\mathbf{X} \sim N(\theta, I_p)$.
- To my knowledge, it was Komaki (2001) who first developed Stein type predictors for the density of a future observation under the normal set up.
- He showed in particular the dominance of the predictive density developed under Stein's (1981) harmonic prior over the classical best equivariant predictive density estimator as developed by Atchison (1975) under the Kullback-Leibler (KL) loss.
- George, Liang and Xu (2006) followed this work and exhibited this dominance under general superharmonic priors.
- They also demonstrated a very nice duality between estimation and prediction under superharmonic priors.

- Some details of the work of George et al. (2006).
- Let $\mathbf{Y} \perp \mathbf{X} | \boldsymbol{\theta}$ where $\mathbf{Y} \sim N(\boldsymbol{\theta}, v_y \mathbf{I}_p)$ and $\mathbf{X} \sim N(\boldsymbol{\theta}, v_x \mathbf{I}_p)$.
- Let $\hat{p}_\pi(\mathbf{y} | \mathbf{X})$ denote the posterior predictive pdf of \mathbf{Y} under an arbitrary prior π .
- Let $m_\pi(\cdot)$ denote a generic notation for marginals.
- Let $W = (v_x^{-1} \mathbf{X} + v_y^{-1} \mathbf{Y}) / (v_x^{-1} + v_y^{-1})$.
- Let $\hat{p}_u(\mathbf{y} | \mathbf{x})$ is the posterior predictive pdf of \mathbf{Y} given \mathbf{X} under the uniform prior $\pi(\boldsymbol{\theta})$; $v_w = (v_x^{-1} + v_y^{-1})^{-1}$.
- Result 1 (The basic identity)

$$\hat{p}_\pi(\mathbf{y} | \mathbf{x}) = [m_\pi(\mathbf{w}; v_w) / m_\pi(\mathbf{x}; v_x)] \hat{p}_u(\mathbf{y} | \mathbf{x})$$
- Result 2. If the prior $\pi(\boldsymbol{\theta})$ is superharmonic, then m_π is superharmonic) or equivalently $m_\pi^{1/2}$ is superharmonic.

- Result 3.

$$R_{KL}(\boldsymbol{\theta}, \hat{p}_u(\mathbf{Y}|\mathbf{X})) - R_{KL}(\boldsymbol{\theta}, \hat{p}_\pi(\mathbf{Y}|\mathbf{X})) = \int_{v_w}^{v_x} \frac{1}{2v^2} \left\{ E_{\boldsymbol{\theta}, v} [\|\mathbf{X} - \boldsymbol{\theta}\|^2 - \|\delta_\pi(\mathbf{X}) - \boldsymbol{\theta}\|^2] \right\} dv,$$

where $\delta_\pi(\mathbf{X})$ is the Bayes estimator of $\boldsymbol{\theta}$ under the prior $\pi(\boldsymbol{\theta})$ relative to the quadratic loss.

- Thus the dominance property of the Bayes point estimator of $\boldsymbol{\theta}$ under squared error loss is inherited by the Bayes predictive densities under KL loss.
- Stein (1981) showed that if $\pi(\boldsymbol{\theta})$ is superharmonic, then $\delta_\pi(\mathbf{X})$ dominates \mathbf{X} under squared error loss.
- By Result 3, the same dominance holds under KL loss.

A Two Sample Problem

- Simultaneous estimation of two population means when one suspects the two means to be nearly equal.
- Rather than estimating the two means separately, it may be more profitable to shrink each mean towards a pooled mean.
- An EB approach handles this in an adaptive way rather than a preliminary test approach.
- Ghosh and Sinha (1988) proposed hierarchical empirical Bayes estimators dominating individual sample means under squared error loss.
- The same phenomenon holds under a general divergence loss.

- \mathbf{X}_1 , \mathbf{X}_2 and S are mutually independent random variables with $\mathbf{X}_i \sim N_p(\boldsymbol{\mu}_i, \sigma^2 \mathbf{V}_i)$, $i = 1, 2$; $S/\sigma^2 \sim \chi_n^2$.
- Notation: $f(\mathbf{x}_1, \mathbf{x}_2 | \boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \sigma^2)$ joint pdf of \mathbf{X}_1 , \mathbf{X}_2 under the above model.
- Divergence loss : $L_\beta(\boldsymbol{\mu}_1, \boldsymbol{\mu}_2; \mathbf{a}_1, \mathbf{a}_2) = \frac{1 - \int f^{1-\beta}(\mathbf{x}_1, \mathbf{x}_2 | \boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \sigma^2) f^\beta(\mathbf{x}_1, \mathbf{x}_2 | \mathbf{a}_1, \mathbf{a}_2, \sigma^2) d\mathbf{x}_1 d\mathbf{x}_2}{\beta(1-\beta)}$.
- g priors: $\boldsymbol{\mu}_i \sim N_p(\boldsymbol{\nu}, \tau^2 \mathbf{V}_i)$, $i = 1, 2$ independently
- $\boldsymbol{\nu} \sim \text{uniform}(R^p)$
- The HB estimators of the $\boldsymbol{\mu}_i$ are given by $\hat{\boldsymbol{\mu}}_i^{HB} = \mathbf{X}_i - B(\mathbf{x}_i - \hat{\boldsymbol{\nu}})$, where $B = \sigma^2/(\sigma^2 + \tau^2)$ and $\hat{\boldsymbol{\nu}} = (\mathbf{V}_1^{-1} + \mathbf{V}_2^{-1})^{-1}(\mathbf{V}_1^{-1} \mathbf{Y}_1 + \mathbf{V}_2^{-1} \mathbf{Y}_2)$.

- First estimate B for known σ^2 .
- Marginally $\mathbf{X}_1 - \mathbf{X}_2 \sim N_p(\mathbf{0}, (\sigma^2/B)(\mathbf{V}_1 + \mathbf{V}_2))$.
- $(\mathbf{X}_1 - \mathbf{X}_2)^T (\mathbf{V}_1 + \mathbf{V}_2)^{-1} (\mathbf{X}_1 - \mathbf{X}_2) / \sigma^2 \sim B^{-1} \chi_p^2$.
- $E(1/\chi_p^2) = (p - 2)^{-1}$ for $p \geq 3$.
- $$\hat{B} = \frac{(p-2)\sigma^2}{(\mathbf{X}_1 - \mathbf{X}_2)^T (\mathbf{V}_1 + \mathbf{V}_2)^{-1} (\mathbf{X}_1 - \mathbf{X}_2)}$$
- Estimate σ^2 by $\hat{\sigma}^2 = S/(n + 2)$.
- $F = (\mathbf{X}_1 - \mathbf{X}_2)^T (\mathbf{V}_1 + \mathbf{V}_2)^{-1} (\mathbf{X}_1 - \mathbf{X}_2) / \hat{\sigma}^2$.

- Consider the shrinkage estimator $\hat{\boldsymbol{\mu}}_i(\phi) = \mathbf{X}_i - \frac{\phi(F)}{F}(\mathbf{X}_i - \hat{\boldsymbol{\nu}})$, $i = 1, 2$.
- Theorem. $(\hat{\boldsymbol{\mu}}_1(\phi), \hat{\boldsymbol{\mu}}_2(\phi))$ dominates $(\mathbf{X}_1, \mathbf{X}_2)$ under the general divergence loss for $p \geq 3$ if
 - (a) $0 < \phi(t) < 2(p - 2)$ for all $t > 0$;
 - (b) $\phi(t)$ is a differentiable nondecreasing function of t for $t > 0$.
- The basic technique for proving this result is to introduce $\mathbf{W} = (\mathbf{V}_1^{-1} + \mathbf{V}_2^{-1})^{-1}(\mathbf{V}_1^{-1}\mathbf{X}_1 + \mathbf{V}_2^{-1}\mathbf{X}_2)$ and exploit the independence of \mathbf{W} with $\mathbf{Z} = \mathbf{X}_1 - \mathbf{X}_2$.

Final Remarks

- Stein's 1956 paper was a major breakthrough in statistics, and its impact is felt even today.
- Stein's original results and many of their variants are continuously used both in the theory and application of statistics.
- The Stein phenomenon for prediction is of more recent origin.
- The present talk has focused on the Stein phenomenon for both estimation and prediction, and has pointed out the duality between the two in several contexts.